

Active Learning of Clustering with Side Information Using ϵ -Smooth Relative Regret Approximations

Nir Ailon and Ron Begleiter

February 1, 2012

Abstract

Clustering is considered a non-supervised learning setting, in which the goal is to partition a collection of data points into disjoint clusters. Often a bound k on the number of clusters is given or assumed by the practitioner. Many versions of this problem have been defined, most notably k -means and k -median.

An underlying problem with the unsupervised nature of clustering is that of determining a similarity function. One approach for alleviating this difficulty is known as clustering with side information, alternatively, semi-supervised clustering. Here, the practitioner incorporates side information in the form of “must be clustered” or “must be separated” labels for data point pairs. Each such piece of information comes at a “query cost” (often involving human response solicitation). The collection of labels is then incorporated in the usual clustering algorithm as either strict or as soft constraints, possibly adding a pairwise constraint penalty function to the chosen clustering objective.

Our work is mostly related to clustering with side information. We ask how to choose the pairs of data points. Our analysis gives rise to a method provably better than simply choosing them uniformly at random. Roughly speaking, we show that the distribution must be biased so as more weight is placed on pairs incident to elements in smaller clusters in some optimal solution. Of course we do not know the optimal solution, hence we don’t know the bias. Using the recently introduced method of ϵ -smooth relative regret approximations of Ailon, Begleiter and Ezra, we can show an iterative process that improves both the clustering and the bias in tandem. The process provably converges to the optimal solution faster (in terms of query cost) than an algorithm selecting pairs uniformly.

1 Introduction

Clustering of data is probably the most important problem in the theory of unsupervised learning. In the most standard setting, the goal is to partition a collection of data points into related groups. Virtually any large scale application using machine learning either uses clustering as a data preprocessing step or as an ends within itself.

In the most traditional sense, clustering is an unsupervised learning problem because the solution is computed from the data itself, with no human labeling involved. There are many versions, most notably k -means and k -median. The number k typically serves as an assumed upper bound on the number of output clusters.

An underlying difficulty with the unsupervised nature of clustering is the fact that a similarity (or distance) function between data points must be chosen by the practitioner as a preliminary

step. This may often not be an easy task. Indeed, even if our dataset is readily embedded in some natural vector (feature) space, we still have the burden of the freedom of choosing a normed metric, and of applying some transformation (linear or otherwise) on the data for good measure. Many approaches have been proposed to tackle this. In one approach, a metric learning algorithm is executed as a preprocessing step in order to choose a suitable metric (from some family). This approach is supervised, and uses distances between pairs of elements as (possibly noisy) labels. The second approach is known as clustering with side information, alternatively, semi-supervised clustering. This approach should be thought of as adding crutches to a lame distance function the practitioner is too lazy to replace. Instead, she incorporates so-called side information in the form of “must be clustered” or “must be separated” labels for data point pairs. Each such label comes at a “query cost” (often involving human response solicitation). The collection of labels is then incorporated in the chosen clustering algorithm as either strict constraints or as soft ones, possibly adding a pairwise constraint penalty function.

1.1 Previous Related Work

Clustering with side information is a fairly new variant of clustering first described, independently, by Demiriz et al. [1999], and Ben-Dor et al. [1999]. In the machine learning community it is also widely known as *semi-supervised clustering*. There are a few alternatives for the form of feedback providing the side-information. The most natural ones are the single item labels [e.g., Demiriz et al., 1999], and the pairwise constraints [e.g., Ben-Dor et al., 1999].

In our study, the side information is pairwise, comes at a cost and is treated frugally. In a related yet different setting, similarity information for all (quadratically many) pairs is available but is noisy. The combinatorial optimization theoretical problem of cleaning the noise is known as *correlation clustering* [Bansal et al., 2002] or *cluster editing* [Shamir et al., 2004]. Constant factor approximations are known for various versions of this problems [Charikar and Wirth, 2004, Ailon et al., 2008]. A PTAS is known for a minimization version in which the number of clusters is fixed [Giotis and Guruswami, 2006].

Roughly speaking, there are two main approaches for utilizing pairwise side information. In the first approach, this information is used to fine tune or learn a *distance* function, which is then passed on to any standard clustering algorithm. Examples include Cohn et al. [2000], Klein et al. [2002], and Xing et al. [2002]. The second approach, which is the starting point to our work, modifies the clustering algorithms’s objective so as to incorporate the pairwise constraints. Basu [2005] in his thesis, which also serves as a comprehensive survey, has championed this approach in conjunction with k -means, and hidden Markov random field clustering algorithms.

1.2 Our Contribution

Our main motivation is reducing the number of pairwise similarity labels (query cost) required for k -clustering data using an *active learning* approach. More precisely, we ask how to choose which pairs of data points to query. Our analysis gives rise to a method provably better than simply choosing them uniformly at random. More precisely, we show that the distribution from which we should draw pairs from must be biased so as more weight is placed on pairs incident to elements in smaller clusters in some optimal solution. Of course we do not know the optimal solution, let alone the bias. Using the recently introduced method of ε -smooth relative regret approximations (ε -SRRA) of Ailon et al. [2011] we can show an iterative process that improves both the clustering

and the bias in tandem. The process provably converges to the optimal solution faster (in terms of query cost) than an algorithm uniformly selecting pairs. Optimality here is with respect to the (complete) pairwise constraint penalty function.

In Section 2 we define our problem mathematically. We then present the ε -SRRA method of Ailon et al. [2011] for the purpose of self containment in Section 3. Finally, we present our main result in Section 4.

2 Notation and Definitions

Let V be a set of points of size n . Our goal is to partition V into k sets (clusters). There are two sources of information guiding us in the process. One is unsupervised, possibly emerging from features attached to each element $v \in V$ together with a chosen distance function. This information is captured in a utility function such as k -means or k -medians. The other type is supervised, and is encoded as an undirected graph $G = (V, E)$. An edge $(u, v) \in E$ corresponds to the constraint *u, v should be clustered together* and a nonedge $(u, v) \notin E$ corresponds to the converse. Each edge or nonedge comes at a *query cost*. This means that G exists only implicitly. We uncover the truth value of the predicate “ $(u, v) \in E$ ” for any chosen pairs u, v for a price. We also assume that G is riddled with human errors, hence it does not necessarily encode a perfect k clustering of the data. In what follows, we assume G fixed.

A k -clustering $\mathcal{C} = \{C_1, \dots, C_k\}$ is a collection of k disjoint (possibly empty) sets satisfying $\bigcup C_i = V$. We use the notation $u \equiv_{\mathcal{C}} v$ if u and v are in the same cluster, and $u \not\equiv_{\mathcal{C}} v$ otherwise.

The cost of \mathcal{C} with respect to G is defined as

$$\text{cost}(\mathcal{C}) = \sum_{(u,v) \in E} \mathbf{1}_{u \not\equiv_{\mathcal{C}} v} + \sum_{(u,v) \notin E} \mathbf{1}_{u \equiv_{\mathcal{C}} v} .$$

Minimizing $\text{cost}(\mathcal{C})$ over clusterings when G is known as correlation clustering (in complete graphs). This problem was defined by Bansal et al. [2004] and has received much attention since (e.g. Ailon et al. [2008], Charikar et al. [2005], Mitra and Samal [2009]).¹ Mitra and Samal [2009] achieved a PTAS for this problem, namely, a polynomial time algorithm returning a k -clustering with cost at most $(1 + \varepsilon)$ that of the optimal.² Their PTAS is not query efficient: It requires knowledge of G in its entirety. In this work we study the query complexity required for achieving a $(1 + \varepsilon)$ approximation for cost. From a learning theoretical perspective, we want to find the best k -clustering explaining G using as few queries as possibly into G .

3 The ε -Smooth Relative Regret Approximation (ε -SRRA) Method

Our search problem can be cast as a special case of the following more general learning problem. Given some possibly noisy structure (e.g. a graph in our case) h , the goal is to find the best explanation using a limited space \mathcal{X} of hypothesis (in our case k -clusterings). The goal is to minimize a notion of a nonnegative *cost* which is defined as the distance $\mathbf{d}(f, h)$ between $f \in \mathcal{X}$ and h . Assume also that the distance function \mathbf{d} between \mathcal{X} and h is an extension of a metric on

¹The original problem definition did not limit the number of output clusters.

²The polynomial degree depends on ε .

- Start with any solution $f_0 \in \mathcal{X}$
- Set $t \leftarrow 0$
- Repeat until some stopping condition:
 - Set $f_t \leftarrow \operatorname{argmin}_{g \in \mathcal{X}} \hat{\Delta}_{f_{t-1}}(g)$, where $\hat{\Delta}_{f_{t-1}}$ is an ε -SRRA for f_{t-1} .
 - Set $t \leftarrow t + 1$

Figure 1: Iterative algorithm using ε -SRRA

\mathcal{X} . Ailon et al. [2011] have recently shown the following general scheme for finding the best $f \in \mathcal{X}$. To explain this scheme, we need to define a notion of ε -smooth relative regret approximation.

Given a solution $f \in \mathcal{X}$ (call it the *pivotal* solution) and another solution $g \in \mathcal{X}$, we define $\Delta_f(g)$ to be $\mathbf{d}(g, h) - \mathbf{d}(f, h)$, namely, the difference between the cost of the solution g and the cost of the solution f . We call this the *relative regret* function with respect to f . Assume we have oracle access to a function $\hat{\Delta}_f : \mathcal{X} \rightarrow \mathbf{R}$ such that for all $g \in \mathcal{X}$,

$$|\hat{\Delta}_f(g) - \Delta_f(g)| \leq \varepsilon \mathbf{d}(f, g) .$$

If such an estimator function $\hat{\Delta}_f$ exists, we say that it is an ε -smooth regret approximation (ε -SRRA) for with respect to f . Ailon et al. [2011] show that if we have an ε -smooth regret approximation function, then it is possible to obtain a $(1 + \varepsilon)$ -approximation to the optimal solution by repeating the iterative process presented in Figure 3. It is shown that this search algorithm converges *exponentially fast* to an $(1 + \varepsilon)$ -approximately optimal one. More precisely, the following is shown:

Theorem 3.1. [Ailon et al., 2011] Assume input parameters $\varepsilon \in (0, 1/5)$ and initializer $\hat{f}_0 \in \mathcal{X}$ of Algorithm 3. Denote $\text{OPT} := \min_{f \in \mathcal{X}} \mathbf{d}(f, h)$. $\hat{f}_0 \in \mathcal{X}$ be an arbitrary function. Then the following holds for \hat{f}_t obtained in Algorithm 3 for all $t \geq 1$:

$$\mathbf{d}(\hat{f}_t, h) \leq (1 + 8\varepsilon) (1 + (5\varepsilon)^t) \text{OPT} + (5\varepsilon)^t \mathbf{d}(\hat{f}_0, h). \quad (3.1)$$

There are two questions now: (1) How can we build $\hat{\Delta}_f$ efficiently? (2) How do we find $\operatorname{argmin}_{g \in \mathcal{X}} \hat{\Delta}_f(g)$?

In the case of k -clusterings, the target structure h is the graph G and \mathcal{X} is the space of k -clusterings over V . The metric \mathbf{d} over \mathcal{X} is taken to be

$$d(\mathcal{C}, \mathcal{C}') = \frac{1}{2} \sum_{u,v} d_{u,v}(\mathcal{C}, \mathcal{C}')$$

where $d_{u,v}(\mathcal{C}, \mathcal{C}') = \mathbf{1}_{u \equiv_{\mathcal{C}'} v} \mathbf{1}_{u \not\equiv_{\mathcal{C}} v} + \mathbf{1}_{u \equiv_{\mathcal{C}} v} \mathbf{1}_{u \not\equiv_{\mathcal{C}'} v}$. By defining $\mathbf{d}(\mathcal{C}, G) := \text{cost}(\mathcal{C})$ we clearly extend \mathbf{d} to a metric over $\mathcal{X} \cup \{G\}$.

4 ε -Smooth Regret Approximation for k -Correlation Clustering

Denote $\text{cost}_{u,v}(\mathcal{C}) = \mathbf{1}_{(u,v) \in E} \mathbf{1}_{u \not\equiv_{\mathcal{C}} v} + \mathbf{1}_{(u,v) \notin E} \mathbf{1}_{u \equiv_{\mathcal{C}} v}$, so that $\text{cost}(\mathcal{C}) = \frac{1}{2} \sum_{u,v} \text{cost}_{u,v}(\mathcal{C})$. Now consider another clustering \mathcal{C}' . We are interested in the change in cost incurred by replacing \mathcal{C} by

\mathcal{C}' , in other words in the function f defined as

$$f(\mathcal{C}') = \text{cost}(\mathcal{C}') - \text{cost}(\mathcal{C}) .$$

We would like to be able to compute an approximation \hat{f} of f by viewing only a sample of edges in G . That is, we imagine that each edge query from G costs us one unit, and we would like to reduce that cost while sacrificing our accuracy as little as possible. We will refer to the cost incurred by queries as the *query complexity*. Consider the following metric on the space of clusterings:

$$d(\mathcal{C}, \mathcal{C}') = \frac{1}{2} \sum_{u,v} d_{u,v}(\mathcal{C}, \mathcal{C}')$$

where $d_{u,v}(\mathcal{C}, \mathcal{C}') = \mathbf{1}_{u \equiv_{\mathcal{C}'} v} \mathbf{1}_{u \not\equiv_{\mathcal{C}} v} + \mathbf{1}_{u \equiv_{\mathcal{C}} v} \mathbf{1}_{u \not\equiv_{\mathcal{C}'} v}$. (The distance function simply counts the number of unordered pairs on which \mathcal{C} and \mathcal{C}' disagree on.) Before we define our sampling scheme, we slightly reorganize the function f . Assume that $|C_1| \geq |C_2| \geq \dots \geq |C_k|$. Denote $|C_i|$ by n_i . The function f will now be written as:

$$f(\mathcal{C}') = \sum_{i=1}^k \sum_{u \in C_i} \left(\frac{1}{2} \sum_{v \in C_i} f_{u,v}(\mathcal{C}') + \sum_{j=i+1}^k \sum_{v \in C_j} f_{u,v}(\mathcal{C}') \right) \quad (4.1)$$

where

$$f_{u,v}(\mathcal{C}') = \text{cost}_{u,v}(\mathcal{C}') - \text{cost}_{u,v}(\mathcal{C}) .$$

Note that $f_{u,v}(\mathcal{C}') \equiv 0$ whenever \mathcal{C} and \mathcal{C}' agree on the pair u, v . For each $i \in [k]$, let $f_i(\mathcal{C}')$ denote the sum running over $u \in C_i$ in (4.1), so that $f(\mathcal{C}') = \sum f_i(\mathcal{C}')$. Similarly, we now rewrite $d(\mathcal{C}, \mathcal{C}')$ as follows:

$$d(\mathcal{C}, \mathcal{C}') = \sum_{i=1}^k \sum_{u \in C_i} \left(\sum_{j=i+1}^k \sum_{v \in C_j} \mathbf{1}_{u \equiv_{\mathcal{C}'} v} + \frac{1}{2} \sum_{v \in C_i} \mathbf{1}_{u \not\equiv_{\mathcal{C}'} v} \right) \quad (4.2)$$

and denote by $d_i(\mathcal{C}')$ the sum over $u \in C_i$ for i fixed in the last expression, so that $d(\mathcal{C}, \mathcal{C}') = \sum_{i=1}^k d_i(\mathcal{C}')$.

Our sampling scheme will be done as follows. Let ε be an error tolerance function, which we set below. For each cluster $C_i \in \mathcal{C}'$ and for each element $u \in C_i$ we will draw $k - i + 1$ independent samples $S_{ui}, S_{u(i+1)}, \dots, S_{uk}$ as follows. Each sample S_{uj} is a subset of C_j of size q (to be defined below), chosen uniformly with repetitions from C_j . We will take

$$q = c_2 k^2 \log n / \varepsilon^4 .$$

where δ is a failure probability (to be used below), and c_2 is a universal constant.

Finally, we define our estimator \hat{f} of f to be:

$$\hat{f}(\mathcal{C}') = \frac{1}{2} \sum_{i=1}^k \frac{|C_i|}{q} \sum_{u \in C_i} \sum_{v \in S_{ui}} f_{u,v}(\mathcal{C}') + \sum_{i=1}^k \sum_{u \in C_i} \sum_{j=i+1}^k \frac{|C_j|}{q} \sum_{v \in S_{uj}} f_{u,v}(\mathcal{C}') .$$

Clearly for each \mathcal{C}' it holds that $\hat{f}(\mathcal{C}')$ is an unbiased estimator of $f(\mathcal{C}')$. We now analyze its error. For each $i, j \in [k]$ let C_{ij} denote $C_i \cap C'_j$. This captures exactly the set of elements in the i 'th cluster in \mathcal{C} and the j 'th cluster in \mathcal{C}' . The distance $d(\mathcal{C}, \mathcal{C}')$ can be written as follows:

$$d(\mathcal{C}, \mathcal{C}') = \frac{1}{2} \sum_{i=1}^k \sum_{j=1}^k |C_{ij} \times (C_i \setminus C_{ij})| + \sum_{j=1}^k \sum_{1 \leq i_1 < i_2 \leq k} |C_{i_1 j} \times C_{i_2 j}|. \quad (4.3)$$

We call each cartesian set product in (4.3) a *distance contributing rectangle*. Note that unless a pair (u, v) appears in one of the distance contributing rectangles, we have $f_{u,v}(\mathcal{C}') = \hat{f}_{u,v}(\mathcal{C}') = 0$. Hence we can decompose \hat{f} and f in correspondence with the distance contributing rectangles, as follows:

$$f(\mathcal{C}') = \frac{1}{2} \sum_{i=1}^k \sum_{j=1}^k F_{i,j}(\mathcal{C}') + \sum_{j=1}^k \sum_{1 \leq i_1 < i_2 \leq k} F_{i_1, i_2, j} \quad (4.4)$$

$$\hat{f}(\mathcal{C}') = \frac{1}{2} \sum_{i=1}^k \sum_{j=1}^k \hat{F}_{i,j}(\mathcal{C}') + \sum_{j=1}^k \sum_{1 \leq i_1 < i_2 \leq k} \hat{F}_{i_1, i_2, j} \quad (4.5)$$

where

$$F_{i,j}(\mathcal{C}') = \sum_{u \in C_{ij}} \sum_{v \in C_i \setminus C_{ij}} f_{u,v}(\mathcal{C}') \quad (4.6)$$

$$\hat{F}_{i,j}(\mathcal{C}') = \frac{|C_i|}{q} \sum_{u \in C_{ij}} \sum_{v \in (C_i \setminus C_{ij}) \cap S_{ui}} f_{u,v}(\mathcal{C}') \quad (4.7)$$

$$F_{i_1, i_2, j}(\mathcal{C}') = \sum_{u \in C_{i_1 j}} \sum_{v \in C_{i_2 j}} f_{u,v}(\mathcal{C}') \quad (4.8)$$

$$\hat{F}_{i_1, i_2, j}(\mathcal{C}') = \frac{|C_{i_2}|}{q} \sum_{u \in C_{i_1 j}} \sum_{v \in C_{i_2 j} \cap S_{ui_2}} f_{u,v}(\mathcal{C}') \quad (4.9)$$

(Note that the S_{ui} 's are multisets, and the inner sums in (4.7) and (4.9) may count elements multiple times.)

Lemma 4.1. *With probability at least $1 - n^{-3}$, the following holds simultaneously for all k -clusterings \mathcal{C}' and all $i, j \in [k]$:*

$$|F_{i,j}(\mathcal{C}') - \hat{F}_{i,j}(\mathcal{C}')| \leq \varepsilon \cdot |C_{ij} \times (C_i \setminus C_{ij})|. \quad (4.10)$$

Proof. Given a k -clustering $\mathcal{C}' = \{C'_1, \dots, C'_k\}$, the predicate (4.10) (for a given i, j) depends only on the set $C_{ij} = C_i \cap C'_j$. Given a subset $B \subseteq C_i$, we say that \mathcal{C}' (i, j) -realizes B if $C_{ij} = B$.

Now fix i, j and $B \subseteq C_i$. Assume a k -clustering (i, j) -realizes B . Let $b = |B|$ and $c = |C_i|$. Consider the random variable $\hat{F}_{i,j}(\mathcal{C}')$ (see (4.7)). Think of the sample S_{ui} as a sequence $S_{ui}(1), \dots, S_{ui}(q)$, where each $S_{ui}(s)$ is chosen uniformly at random from C_i for $s = 1, \dots, q$. We can now rewrite $\hat{F}_{i,j}(\mathcal{C}')$ as follows:

$$\hat{F}_{i,j}(\mathcal{C}') = \frac{c}{q} \sum_{u \in B} \sum_{s=1}^q X(S_{ui}(s))$$

where

$$X(v) = \begin{cases} f_{u,v}(\mathcal{C}') & v \in C_i \setminus C_{ij} \\ 0 & \text{otherwise} \end{cases}.$$

For all $s = 1, \dots, q$ the random variable $X(S_{ui}(s))$ is bounded by 2 almost surely, and its moments satisfy:

$$\begin{aligned} E[X(S_{ui}(s))] &= \frac{1}{c} \sum_{v \in (C_i \setminus C_{ij})} f_{u,v}(\mathcal{C}') \\ E[X(S_{ui}(s))^2] &\leq \frac{4(c-b)}{c}. \end{aligned} \quad (4.11)$$

From this we conclude using Bernstein inequality that for any $t \leq b(c-b)$,

$$\Pr[|\hat{F}_{i,j}(\mathcal{C}') - F_{i,j}(\mathcal{C}')| \geq t] \leq \exp \left\{ -\frac{qt^2}{16cb(c-b)} \right\}$$

Plugging in $t = \varepsilon b(c-b)$, we conclude

$$\Pr[|\hat{F}_{i,j}(\mathcal{C}') - F_{i,j}(\mathcal{C}')| \geq \varepsilon b(c-b)] \leq \exp \left\{ -\frac{q\varepsilon^2 b(c-b)}{16c} \right\}$$

Now note that the number of possible sets $B \subseteq C_i$ of size b is at most $n^{\min\{b, c-b\}}$. Using union bound and recalling our choice of q , the lemma follows. \square

The Lemma can be easily proven using the Bernstein probability inequality. A bit more involved is the following:

Lemma 4.2. *With probability at least $1 - n^{-3}$, the following holds simultaneously for all k -clusterings \mathcal{C}' and for all $i_1, i_2, j \in [k]$ with $i_1 < i_2$:*

$$|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| \leq \varepsilon \max \left\{ |C_{i_1 j} \times C_{i_2 j}|, \frac{|C_{i_1 j} \times (C_{i_1} \setminus C_{i_1 j})|}{k}, \frac{|C_{i_2 j} \times (C_{i_2} \setminus C_{i_2 j})|}{k} \right\} \quad (4.12)$$

Proof. Given a k -clustering $\mathcal{C}' = \{C'_1, \dots, C'_k\}$, the predicate (4.12) (for a given i_1, i_2, j) depends only on the sets $C_{i_1 j} = C_{i_1} \cap C'_j$ and $C_{i_2 j} = C_{i_2} \cap C'_j$. Given subsets $B_1 \subseteq C_{i_2}$ and $B_2 \subseteq C_{i_2}$, we say that \mathcal{C}' (i_1, i_2, j) -realizes (B_1, B_2) if $C_{i_1 j} = B_1$ and $C_{i_2 j} = B_2$.

We now fix $i_1 < i_2, j$ and $B_1 \subseteq C_{i_1}, B_2 \subseteq C_{i_2}$. Assume a k -clustering \mathcal{C}' (i_1, i_2, j) -realizes (B_1, B_2) . For brevity, denote $b_\iota = |B_\iota|$ and $c_\iota = |C_{i_\iota}|$ for $\iota = 1, 2$. Using Bernstein inequality as before, we conclude that

$$\Pr[|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| > t] \leq \exp \left\{ -\frac{c_3 t^2 q}{b_1 b_2 c_2} \right\}. \quad (4.13)$$

for any t in the range $[0, b_1 b_2]$, for some global $c_4 > 0$. For t in the range $(b_1 b_2, \infty)$,

$$\Pr[|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| > t] \leq \exp \left\{ -\frac{c_5 t q}{c_2} \right\}. \quad (4.14)$$

We consider the following three cases.

1. $b_1 b_2 \geq \max\{b_1(c_1 - b_1/k, b_2(c_2 - b_2)/k)\}$. Hence, $b_1 \geq (c_2 - b_2)/k, b_2 \geq (c_1 - b_1)/k$. In this case, plugging in (4.13) we get

$$\Pr[|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| > \varepsilon b_1 b_2] \leq \exp \left\{ -\frac{c_3 \varepsilon^2 b_1 b_2 q}{c_2} \right\}. \quad (4.15)$$

Consider two subcases. (i) If $b_2 \geq c_2/2$ then the RHS of (4.15) is at most $\exp \left\{ -\frac{c_3 \varepsilon^2 b_1 q}{2} \right\}$. The number of sets B_1, B_2 of sizes b_1, b_2 respectively is clearly at most $n^{b_1 + (c_2 - b_2)} \leq n^{b_1 + k b_1}$. Therefore, if $q = O(\varepsilon^{-2} k \log n)$, then with probability at least $1 - n^{-6}$ simultaneously for all B_1, B_2 of sizes b_1, b_2 respectively and for all \mathcal{C}' (i_1, i_2, j) -realizing (B_1, B_2) we have that $|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| \leq \varepsilon b_1 b_2$. (ii) If $b_2 < c_2/2$ then by our assumption, $b_1 \geq c_2/2k$. Hence the RHS of (4.15) is at most $\exp \left\{ -\frac{c_3 \varepsilon^2 b_2 q}{2k} \right\}$. The number of sets B_1, B_2 of sizes b_1, b_2 respectively is clearly at most $n^{(c_1 - b_1) + c_2} \leq n^{b_2(1+k)}$. Therefore, if $q = O(\varepsilon^{-2} k^2 \log n)$, then with probability at least $1 - n^{-6}$ simultaneously for all B_1, B_2 of sizes b_1, b_2 respectively and for all \mathcal{C}' (i_1, i_2, j) -realizing (B_1, B_2) we have that $|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| \leq \varepsilon b_1 b_2$.

2. $b_2(c_2 - b_2)/k \geq \max\{b_1 b_2, b_1(c_1 - b_1)/k\}$. We consider two subcases.

(a) $\varepsilon b_2(c_2 - b_2)/k \leq b_1 b_2$. Using (4.13), we get

$$\Pr[|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| > \varepsilon b_2(c_2 - b_2)/k] \leq \exp \left\{ -\frac{c_3 \varepsilon^2 b_2(c_2 - b_2)^2 q}{k^2 b_1 c_2} \right\} \quad (4.16)$$

Again consider two subcases. (i) $b_2 \leq c_2/2$. In this case we conclude from (4.16)

$$\Pr[|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| > \varepsilon b_2(c_2 - b_2)/k] \leq \exp \left\{ -\frac{c_3 \varepsilon^2 b_2 c_2 q}{4k^2 b_1} \right\} \quad (4.17)$$

Now note that by assumption

$$b_1 \leq (c_2 - b_2)/k \leq c_2/k \leq c_1/k. \quad (4.18)$$

Also by assumption, $b_1 \leq b_2(c_2 - b_2)/(c_1 - b_1) \leq b_2 c_2/(c_1 - b_1)$. Plugging in (4.18), we conclude that $b_1 \leq b_2 c_2/(c_1(1 - 1/k)) \leq 2b_2 c_2/c_1 \leq 2b_2$. From here we conclude that the RHS of (4.17) is at most $\exp \left\{ -\frac{c_3 \varepsilon^2 c_2 q}{4k^2} \right\}$. The number of sets B_1, B_2 of sizes b_1, b_2 respectively is clearly at most $n^{b_1 + b_2} \leq n^{2b_2 + b_2} \leq n^{3c_2}$. Hence, if $q = O(\varepsilon^{-2} k^2 \log n)$ then with probability at least $1 - n^{-6}$ simultaneously for all B_1, B_2 of sizes b_1, b_2 respectively and for all \mathcal{C}' (i_1, i_2, j) -realizing (B_1, B_2) we have that $|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| \leq \varepsilon b_2(c_2 - b_2)/k$. In the second subcase (ii) $b_2 > c_2/2$. The RHS of (4.16) is at most $\exp \left\{ -\frac{c_3 \varepsilon^2 (c_2 - b_2)^2 q}{2k^2 b_1} \right\}$. By our assumption, $(c_2 - b_2)/(k b_1) \geq 1$, hence this is at most $\exp \left\{ -\frac{c_3 \varepsilon^2 (c_2 - b_2) q}{2k} \right\}$. The number of sets B_1, B_2 of sizes b_1, b_2 respectively is clearly at most $n^{b_1 + (c_2 - b_2)} \leq n^{(c_2 - b_2)/k + (c_2 - b_2)} \leq n^{2(c_2 - b_2)}$. Therefore, if $q = O(\varepsilon^{-2} k \log n)$, then with probability at least $1 - n^{-6}$ simultaneously for all B_1, B_2 of sizes b_1, b_2 respectively and for all \mathcal{C}' (i_1, i_2, j) -realizing (B_1, B_2) we have that $|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| \leq \varepsilon b_2(c_2 - b_2)/k$.

(b) $\varepsilon b_2(c_2 - b_2)/k > b_1 b_2$. We now use (4.14) to conclude

$$\Pr[|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| > \varepsilon b_2(c_2 - b_2)/k] \leq \exp \left\{ -\frac{c_5 \varepsilon b_2(c_2 - b_2)q}{k c_2} \right\} \quad (4.19)$$

We again consider the cases (i) $b_2 \leq c_2/2$ and (ii) $b_2 \geq c_2/2$ as above. In (i), we get that the RHS of (4.19) is at most $\exp \left\{ -\frac{c_5 \varepsilon b_2 q}{2k} \right\}$, that $b_1 \leq 2b_2$ and hence the number of possibilities for B_1, B_2 is at most $n^{b_1+b_2} \leq n^{3b_2}$. In (ii), we get that the RHS of (4.19) is at most $\exp \left\{ -\frac{c_5 \varepsilon (c_2 - b_2)q}{2k} \right\}$, and the number of possibilities for B_1, B_2 is at most $n^{2(c_2 - b_2)}$. For both (i) and (ii) taking $q = O(\varepsilon^{-1} k \log n)$ ensures with probability at least $1 - n^{-6}$ simultaneously for all B_1, B_2 of sizes b_1, b_2 respectively and for all \mathcal{C}' (i_1, i_2, j) -realizing (B_1, B_2) we have that $|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| \leq \varepsilon b_2(c_2 - b_2)/k$.

3. $b_1(c_1 - b_1)/k \geq \max\{b_1 b_2, b_2(c_2 - b_2)/k\}$. We consider two subcases.

- $\varepsilon b_1(c_1 - b_1)/k \leq b_1 b_2$. Using (4.13), we get

$$\Pr[|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| > \varepsilon b_1(c_1 - b_1)/k] \leq \exp \left\{ -\frac{c_3 \varepsilon^2 b_1(c_1 - b_1)^2 q}{k^2 b_2 c_2} \right\}. \quad (4.20)$$

As before, consider case (i) in which $b_2 \leq c_2/2$ and (ii) in which $b_2 \geq c_2/2$. For case (i), we notice that the RHS of (4.19) is at most $\exp \left\{ -\frac{c_3 \varepsilon^2 b_2(c_2 - b_2)(c_1 - b_1)q}{k^2 b_2 c_2} \right\}$ (we used the fact that $b_1(c_1 - b_1) \geq b_2(c_2 - b_2)$ by assumption). This is hence at most $\exp \left\{ -\frac{c_3 \varepsilon^2 (c_1 - b_1)q}{2k^2} \right\}$. The number of possibilities of B_1, B_2 of sizes b_1, b_2 is clearly at most $n^{(c_1 - b_1) + b_2} \leq n^{(c_1 - b_1) + (c_1 - b_1)/k} \leq n^{2(c_1 - b_1)}$. From this we conclude that $q = O(\varepsilon^{-2} k^2 \log n)$ suffices for this case. For case (ii), we bound the RHS of (4.20) by $\exp \left\{ -\frac{c_3 \varepsilon^2 b_1(c_1 - b_1)^2 q}{2k^2 b_2^2} \right\}$. Using the assumption that $(c_1 - b_1)/b_2 \geq k$, the latter expression is upper bounded by $\exp \left\{ -\frac{c_3 \varepsilon^2 b_1 q}{2} \right\}$. Again by our assumptions,

$$b_1 \geq b_2(c_2 - b_2)/(c_1 - b_1) \geq (\varepsilon(c_1 - b_1)/k)(c_2 - b_2)/(c_1 - b_1) = \varepsilon(c_2 - b_2)/k. \quad (4.21)$$

The number of possibilities of B_1, B_2 of sizes b_1, b_2 is clearly at most $n^{b_1 + (c_2 - b_2)}$ which by (4.21) is bounded by $n^{b_1 + kb_1/\varepsilon} \leq n^{2kb_1/\varepsilon}$. From this we conclude that $q = O(\varepsilon^{-3} k \log n)$ suffices for this case.

- $\varepsilon b_1(c_1 - b_1)/k > b_1 b_2$.

$$\Pr[|F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| > \varepsilon b_1(c_2 - b_1)/k] \leq \exp \left\{ -\frac{c_5 \varepsilon b_1(c_1 - b_1)q}{k c_2} \right\} \quad (4.22)$$

We consider two sub-cases, (i) $b_1 \leq c_1/2$ and (ii) $b_1 > c_1/2$. In case (i), we have that

$$\begin{aligned} \frac{b_1(c_1 - b_1)}{c_2} &= \frac{1}{2} \frac{b_1(c_1 - b_1)}{c_2} + \frac{1}{2} \frac{b_1(c_1 - b_1)}{c_2} \\ &\geq \frac{1}{2} \frac{b_1 c_1}{2c_2} + \frac{1}{2} \frac{b_2(c_2 - b_2)}{c_2} \\ &\geq b_1/4 + \min\{b_2, c_2 - b_2\}/2. \end{aligned}$$

Hence, the RHS of (4.22) is bounded above by $\exp \left\{ -\frac{c_5 \varepsilon q (b_1/4 + \min\{b_2, c_2 - b_2\}/2)}{k} \right\}$. The number of possibilities of B_1, B_2 of sizes b_1, b_2 is clearly at most $n^{b_1 + \min\{b_2, c_2 - b_2\}}$, hence it suffices to take $q = O(\varepsilon^{-1} k \log n)$ for this case. In case (ii), we can upper bound the RHS of (4.22) by $\exp \left\{ -\frac{c_5 \varepsilon c_1 (c_1 - b_1) q}{2k c_2} \right\} \geq \exp \left\{ -\frac{c_5 \varepsilon (c_1 - b_1) q}{2k} \right\}$. The number of possibilities of B_1, B_2 of sizes b_1, b_2 is clearly at most $n^{(c_1 - b_1) + b_2}$ which, using our assumptions, is bounded above by $n^{(c_1 - b_1) + (c_1 - b_1)/k} \leq n^{2(c_1 - b_2)}$. Hence, it suffices to take $q = O(\varepsilon^{-1} k \log n)$ for this case.

This concludes the proof of the lemma. \square

As a consequence, we get the following:

Lemma 4.3. *with probability at least $1 - n^{-3}$, the following holds simultaneously for all k -clusterings \mathcal{C}' :*

$$|f(\mathcal{C}') - f(\mathcal{C})| \leq 3\varepsilon d(\mathcal{C}', \mathcal{C}) .$$

Proof.

$$\begin{aligned}
|f(\mathcal{C}') - \hat{f}(\mathcal{C}')| &= \frac{1}{2} \sum_{i=1}^k \sum_{j=1}^k |F_{i,j}(\mathcal{C}') - \hat{F}_{i,j}(\mathcal{C}')| + \sum_{j=1}^k \sum_{1 \leq i_1 < i_2 \leq k} |F_{i_1, i_2, j}(\mathcal{C}') - \hat{F}_{i_1, i_2, j}(\mathcal{C}')| \\
&\leq \frac{1}{2} \sum_{i=1}^k \sum_{j=1}^k \varepsilon^{-2} k^{-1} |C_{ij} \times (C_i \setminus C_{ij})| \\
&\quad + \varepsilon \sum_{j=1}^k \sum_{i_1 < i_2} (|C_{i_1 j} \times C_{i_2 j}| + k^{-1} |C_{i_1 j} \times (C_{i_1} \setminus C_{i_1 j})| + k^{-1} |C_{i_2 j} \times (C_{i_2} \setminus C_{i_2 j})|) \\
&\leq \frac{1}{2} \sum_{i=1}^k \sum_{j=1}^k \varepsilon^{-2} k^{-1} |C_{ij} \times (C_i \setminus C_{ij})| + \varepsilon \sum_{j=1}^k \sum_{i_1 < i_2} |C_{i_1 j} \times C_{i_2 j}| \\
&\quad + \varepsilon \sum_{j=1}^k \sum_{i_1=1}^k \sum_{i_2=1}^k k^{-1} |C_{i_1 j} \times (C_{i_1} \setminus C_{i_1 j})| + \varepsilon \sum_{j=1}^k \sum_{i_2=1}^k \sum_{i_1=1}^k k^{-1} |C_{i_2 j} \times (C_{i_2} \setminus C_{i_2 j})| \\
&\leq \frac{1}{2} \sum_{i=1}^k \sum_{j=1}^k \varepsilon^{-2} k^{-1} |C_{ij} \times (C_i \setminus C_{ij})| + \varepsilon \sum_{j=1}^k \sum_{i_1 < i_2} |C_{i_1 j} \times C_{i_2 j}| \\
&\quad + \varepsilon \sum_{j=1}^k \sum_{i_1=1}^k k k^{-1} |C_{i_1 j} \times (C_{i_1} \setminus C_{i_1 j})| + \varepsilon \sum_{j=1}^k \sum_{i_2=1}^k k k^{-1} |C_{i_2 j} \times (C_{i_2} \setminus C_{i_2 j})| \\
&\leq \varepsilon \frac{3}{2} \sum_{i=1}^k \sum_{j=1}^k |C_{ij} \times (C_i \setminus C_{ij})| + \varepsilon \sum_{j=1}^k \sum_{i_1 < i_2} |C_{i_1 j} \times C_{i_2 j}| \\
&\leq 3\varepsilon d(\mathcal{C}, \mathcal{C}')
\end{aligned}$$

The first equality was (4.4)-(4.5). The second was Lemmas 4.1- 4.2 (assuming success of a high probability event), the third, fourth and fifth inequalities were rearrangement of the sum, and the final inequality came from (4.3). \square

5 Conclusions and Future Work

Our study considered the information theoretical problem of choosing which questions to ask in a game in which adversarially noisy combinatorial pairwise information is input to a clustering algorithm. We designed and analyzed a distribution from which drawing pairs is provably superior than the uniform distribution. Our analysis did not take into account geometric information (e.g. a feature vector attached to each data point) and treated the similarity labels as side information, as suggested in a recent line of literature. It would be interesting to study our solution in conjunction with geometric information. It would also be interesting to study our approach in the context of metric learning, where the goal is to cleverly choose which pairs to obtain (noisy) distance labels for.

References

- Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating inconsistent information: Ranking and clustering. *J. ACM*, 55(5):1–27, 2008.
- Nir Ailon, Ron Begleiter, and Esther Ezra. A new active learning scheme with applications to learning to rank from pairwise preferences. In *arXiv:1110.2136*, 2011.
- Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. In *MACHINE LEARNING*, pages 238–247, 2002.
- Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. *Machine Learning*, 56: 89–113, 2004.
- Sugato Basu. *Semi-supervised Clustering: Probabilistic Models, Algorithms and Experiments*. PhD thesis, Department of Computer Sciences, University of Texas at Austin, 2005.
- Amir Ben-Dor, Ron Shamir, and Zohar Yakhini. Clustering gene expression patterns. *Journal of Computational Biology*, 6(3/4):281–297, 1999.
- Moses Charikar and Anthony Wirth. Maximizing quadratic programs: Extending grothendieck’s inequality. In *FOCS*, pages 54–60. IEEE Computer Society, 2004.
- Moses Charikar, Venkatesan Guruswami, and Anthony Wirth. Clustering with qualitative information. *J. Comput. Syst. Sci.*, 71(3):360–383, 2005.
- David Cohn, Rich Caruana, and Andrew McCallum. Semi-supervised clustering with user feedback. unpublished manuscript, 2000. URL <http://www.cs.umass.edu/~mccallum/papers/semisup-aaai2000s.ps>.
- Ayhan Demiriz, Kristin Bennett, and Mark J. Embrechts. Semi-supervised clustering using genetic algorithms. In *In Artificial Neural Networks in Engineering (ANNIE-99)*, pages 809–814. ASME Press, 1999.
- Ioannis Giotis and Venkatesan Guruswami. Correlation clustering with a fixed number of clusters. *Theory of Computing*, 2(1):249–266, 2006.

- Dan Klein, Sepandar D. Kamvar, and Christopher D. Manning. From instance-level constraints to space-level constraints: Making the most of prior knowledge in data clustering. In *ICML*, pages 307–314, 2002.
- P. Mitra and M. Samal. Approximation algorithm for correlation clustering. In *Networked Digital Technologies, 2009. NDT '09. First International Conference on*, pages 140 –145, july 2009.
- Ron Shamir, Roded Sharan, and Dekel Tsur. Cluster graph modification problems. *Discrete Applied Math*, 144:173–182, November 2004.
- Eric P. Xing, Andrew Y. Ng, Michael I. Jordan, and Stuart Russell. Distance metric learning, with application to clustering with side-information. In *Advances in Neural Information Processing Systems 15*, pages 505–512. MIT Press, 2002.